arXiv:1702.05839v1 [cs.CV] 20 Feb 2017

Progressively Diffused Networks for Semantic Image Segmentation

Ruimao Zhang [1]*, Wei Yang [2]*, Zhanglin Peng [3], Xiaogang Wang [2], Liang Lin [1]

[1] Sun Yat-sen University, Guangzhou, China
[2] The Chinese University of Hong Kong, Hong Kong, China
[3] SenseTime Group Limited, Shenzhen, China

Abstract

This paper introduces Progressively Diffused Networks (PDNs) for unifying multi-scale context modeling with deep feature learning, by taking semantic image segmentation as an exemplar application. Prior neural networks such as ResNet [11] tend to enhance representational power by increasing the depth of architectures and driving the training objective across layers. However, we argue that spatial dependencies in different layers, which generally represent the rich contexts among data elements, are also critical to building deep and discriminative representations. To this end, our PDNs enables to progressively broadcast information over the learned feature maps by inserting a stack of information diffusion layers, each of which exploits multi-dimensional convolutional LSTMs (Long-Short-Term Memory Structures). In each LSTM unit, a special type of atrous filters are designed to capture the short range and long range dependencies from various neighbors to a certain site of the feature map and pass the accumulated information to the next layer. From the extensive experiments on semantic image segmentation benchmarks (e.g., ImageNet Parsing, PASCAL VOC2012 and PASCAL-Part), our framework demonstrates the effectiveness to substantially improve the performances over the popular existing neural network models, and achieves state-of-the-art on ImageNet Parsing for large scale semantic segmentation.

1. Introduction

In the literature, representation learning [37] is a set of methods to automatically discover the representation of raw data for intelligent tasks. Recent developed Deep Convolutional Neural Networks (CNN) [17, 16, 26, 22, 11] is one of representation learning method with multiple levels of representation, which transforms the raw input image into abstract features by stacking several non-linear modules. To enhance the representation power of such architecture, prior works [26, 29, 27, 11] focus on increasing the depth of architectures and driving the training objective across layers. These methods achieve great success on the challenge ImageNet competition [16, 11], and the semantic segmentation task has also greatly benefited from such "very deep" model [4]. However, increasing network depth makes it difficult to design the architecture, while the gradient vanishing makes the model hard to train.

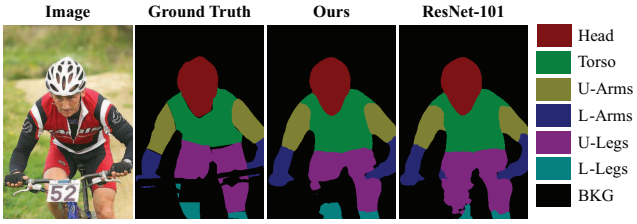

Figure 1. An example of semantic image segmentation with and without incorporating spatial dependencies. From left to right are input image, groundtruth labeling, segmentation result by proposed PDNs and segmentation result by ResNet-101 [11]. It is obvious that our architecture can capture the rich contexts among pixels and achieve the more reasonable global reasoning.

On the other hand, another branch of works trying to explore rich contexts in the visual data have also received much attention in literature [15, 32, 35]. Recent evidence reveals incorporating graphic model, e.g. CRF [4] or MRF [21, 24] to the output confidence maps of CNN is of crucial importance, which greatly improves the accuracy of dense prediction. But these context modeling methods require careful pairwise constraints design and do not explicitly enhance the pixel-wise representation, leading suboptimal segmentation results.

An alternative scheme focuses on exploiting Long Short Term Memory (LSTM) networks to automatically learn the spatial dependencies. These data-driven methods use contextual information to enhance intermediate feature representations, achieving promising results on semantic segmentation task [2, 30, 19, 18], where the property of long-range dependencies is used to pass the information between

*The first two authors contribute equally to this work.

neighbor pixels layer by layer. However, in terms of information diffusion range in each layer, most existing approaches [14, 31, 19, 18] have only explored well-designed short distance. A wider range of information diffusion is achieved by stacking multiple LSTM layers. As illustrated in Fig. 3 (a) and Fig. 3 (b), the feature enhancement of each position by above methods is determined by the short distance neighbors (e.g. the closest 2 to 8 adjacent positions) in each layer, limiting the breadth and speed of information diffusion. At the same time, all of the existing methods are based on fully-connected LSTM, and the computational cost of these methods is another obvious limitation.

In this paper, we propose a novel Progressively Diffused Networks (PDNs) that extends the traditional neural network structure from learning rather complex pattern layer by layer to spreading the contextual information in image plane, and demonstrate its superiority on various semantic segmentation tasks. PDNs introduces a stack of information diffusion layers for context modeling, each of which contains several multi-dimensional Long Short-Term Memory (LSTM) networks. Instead of spreading information to a fixed number of adjacent positions, the diffusion layer allows to propagate information from a certain position to a large range in the image plane. It can effectively expand the scope of communication in a single layer, while improving the speed of information diffusion in multi layers.

Specifically, we propose two types of diffused LSTM, one called spatial LSTM and another called depth LSTM. Each spatial LSTM in the diffusion layer generates a certain type of contextual feature maps. Intuitively, these contextual feature maps have different meanings with the ones generated by traditional convolutional methods. The value in the convolutional maps represents the response of a local area under a certain pattern. In contrast, each site of the contextual feature maps involves the information it will propagate to its neighbors in the next state. Different from spatial LSTM, we incorporate the depth LSTM [19] into diffusion layer to realize the communication of each site from one layer to the next.

In each diffused LSTM unit, the special atrous filters [5] for each contextual feature maps are used for each position to capture the diverse neighborhood information in a large range of local area. Finally, these filtering results will be integrated to calculate the information of each site passed to its neighbors or to itself in the next layer. Compared with the fully-connected LSTM in previous works [14, 30, 19, 18], this convolution-based version is more intuitive, and can significantly improve the computational efficiency.

This paper has following four contributions. (1) We propose a novel deep architecture named Progressively Diffused Networks. The stacked diffusion layers allow the contextual information spread from a certain position to a large range on natural image. (2) We design a novel LSTM-based layer, which contains several convolutional LSTM. It can generate a series of contextual feature maps, and the convolutional LSTM in the next layer can exploit these maps to further guide the feature representation of each site. (3) A special type of atrous filters are incorporated into proposed convolutional LSTM, each of which is corresponding to a special contextual feature map. Through the convolutional operation, we reduce the parameter space in context modeling. (4) We obtain state-of-art results on three challenging datasets, i.e. ImageNet Parsing Dataset [39], PASCAL-Part Dataset [7] and PASCAL VOC 2012 semantic segmentation benchmark [8].

2. Related Work

The performance of computer vision tasks is heavily dependent on the choice of visual representation. For that reason, many of previous efforts in deploying computer vision models focused on designing the pipelines to extract the effective visual representation [36, 9, 23]. Such feature engineering based methods are important but require a lot of domain knowledge, severely limiting the development of visual applications. In order to make the vision models less dependent on feature engineer, representation learning [1, 37], which facilitates useful information extraction from raw data for building predictors, has attracted much attention in the past decade.

A typical representation learning model is Convolutional Neural Networks (CNN) [17, 16, 22], which is designed to process the data with multiple arrays such as images [16] or videos [12]. By stacking several convolution-pooling layers, this model transforms the visual representation from one level into a slightly more abstract level. Recently, many works tend to enhance representational power of CNN by increasing the depth of architectures [26, 29, 27, 11], and achieve great success on image classification [16, 11]. The dense prediction task, such as semantic segmentation, has also benefited from such deep feature learning methods [22, 4]. In [22], Long *et al*. firstly replaced fully-connected layers of CNN with convolutional layers, making it possible to accomplish pixel-wise prediction in the whole image by the deep model. Chen *et al*. [5] further proposed the atrous convolution to explicitly control the resolution of feature responses, and exhibited the atrous spatial pyramid pooling for dense predicting at multiple scales.

Meanwhile, in order to explicitly discover the intricate structures in the visual data for dense labeling, the graphic models were applied to explore the rich information (e.g. long-range dependencies or high-order potentials) in the image by defining the spatial constrains. In [4], the confidence maps generated by the Fully Convolutional Networks (FCN) [22] were fed into the Conditional Random Field (CRF) with simple pairwise potentials for post-processing, but this model treated the FCN and CRF as separated com-

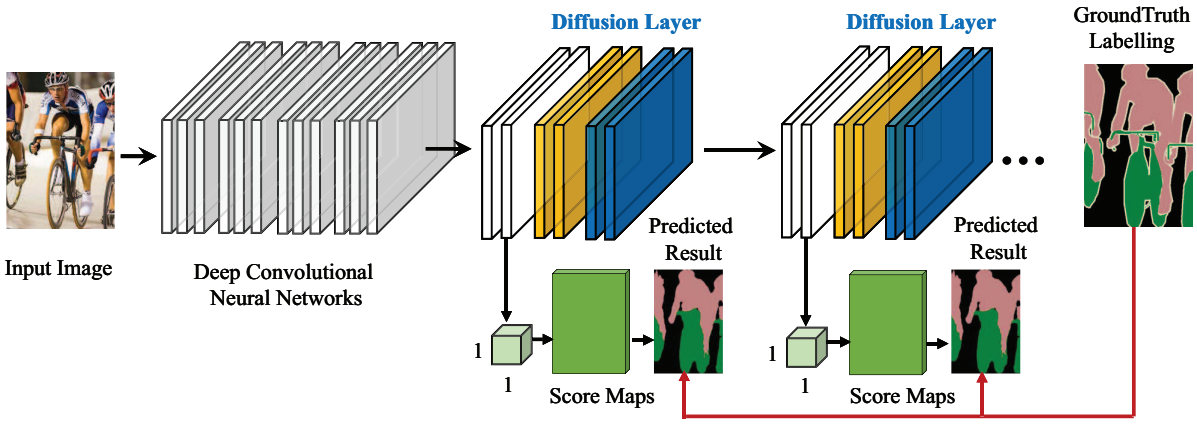

Figure 2. The framework of proposed Progressively Diffused Networks for semantic segmentation. The diffusion module with several diffusion layers is stacked on the top of Deep Convolutional Neural Networks for context modeling. Each diffusion layer outputs several contextual feature maps (i.e. yellow and blue maps) for broadcasting neighborhood information on the image plane in the next layer. And depth feature maps (i.e. white maps) are also generated to communicate information of each site from one layer to the next. For the model training, the depth feature maps in each layer are convolved with 1×1 convolutional filters to generate the score maps for cost calculation. In the test phase, all of the score maps from distinct diffusion layers will be merged to predict the final result.

ponents, limiting the joint optimization of the model. In contrast, Schwing *et al*. [24] jointly train the FCN and Markov Random Field (MRF) by passing the error generated by MRF back to the neural networks. However, the iterative inference algorithm (i.e. Mean Field inference) used in this method is time consuming. To improve computational efficiency, Liu *et al*. [21] solve MRF by the convolution operations, which devises the additional layers to approximate the mean field inference for pairwise terms. Although these methods significantly improve the performance of dense labelling, the contextual information is still not explicitly encoded into the pixel-wise representations.

In the literature, the Long Short Term Memory (LSTM) Network has been introduced to deal with the long-range dependencies in the representation modeling, and this advanced Recurrent Neural Network (RNN) has achieved great success in many intelligent tasks [10, 28, 34, 25]. In recent years, it has been extended to multi-dimensional communication [14, 30, 40] and adapted to represent the rich contexts in image spatial [19, 18]. In [19], a recent advance in LSTM-based context modeling was achieved by considering both short dependencies from local area and long-distance global information from the whole image. Liang *et al*. [18] further extended this work from multi-dimensional data to general graph-structured data, and constructed an adaptive graph topology to propagate contextual information between adjacent superpixels. Nevertheless, in these works, the feature representation of each position is affected by a limited local factors (i.e. the adjacent positions), which restricts the capacity of involving diverse visual correlations in a large range. Different from using limited local LSTM units, the proposed PDNs captures the short-range and long-range dependencies from various neighbors and can generate more informative representation for pixel-wise prediction.

3. Progressively Diffused Network

In this section, we will first review the framework of proposed PDNs and introduce the diffusion mechanism for context modeling. The implementation details of diffusion layer will be described at the end of section.

3.1. Framework Overview

We develop a novel Progressively Diffused Networks (PDNs) for image context modeling, and pursue the semantic segmentation task. The architecture of proposed network is presented in Fig. 2. The input image is first forwarded through the Deep Convolutional Neural Networks (e.g. ResNet-101 [11]) to generate a set of feature maps. Then these feature maps are fed into a series of diffusion layers to progressively spread the context information on the image plane. After each diffused LSTM layer, the generated depth maps (i.e. white maps in Fig. 2) are convolved with 1×1 filters to calculate the score maps for dense prediction. For the model training, the cross-entropy loss over all pixels is used for every diffused LSTM layer. In the test phase, the final result is predicted by merging all of the score maps from the distinct diffusion layers.

3.2. Diffusion Mechanism

The diffusion mechanism aims at broadcasting the contextual information on the image plane and exploiting such information to increase the discrimination of feature representation for each pixel. In fact, such mechanism is partially inspired by biological research that the pheromone released by a single cell can affect larger tissue areas, not just the adjacent cells [3]. In this article, the diffusion mechanism is

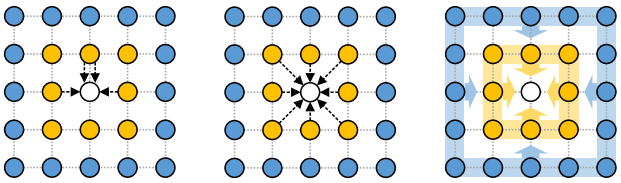

(a) Diagonal BiLSTM (b) Local-Global LSTM (c) Diffused LSTM

Figure 3. The comparison of different types of LSTM unit. Subfigure (a) and (b) show the previous pixel-wise LSTM unit (i.e. Diagonal BiLSTM [31] and LG-LSTM[19]) that update the states of each site by adopting fixed local factors (i.e. adjacent sites). In (c), our proposed LSTM unit can capture the short-range and long-range dependencies from the divers neighbors and can generate more informative data representation

adopted to spread information from one site to a large field of neighbors. It includes two coherent aspects, (1) using the information from different neighbors to enrich the feature representation of one certain site, illustrated in Fig. 3; (2) propagating different information from the certain site to its different neighbors to guide their further representation, illustrated in Fig. 4.

The diffusion mechanism in this paper is similar to recent image processing works based on LSTM [14, 30, 31, 19]. However, in these works, the contextual information of each position is determined by a fixed factorization (e.g. 2 to 8 neighboring positions), as shown in Fig. 3 (a) and Fig. 3 (b). Different from these locally fixed LSTM units, the modified LSTM in our PDNs allows each location to receive messages from different numbers of neighbors, as illustrated in Fig. 3 (c). For most of the previous approaches [14, 19], the parameters of each LSTM are shared, thus the information that each site passes to all of its neighbors is equivalent. Therefore, these methods can be viewed as a special case of proposed diffusion mechanism.

3.3. Diffusion Layer

In practise, the diffusion layer exploits the multi-dimensional convolutional LSTMs to receive and broadcast the information. With these convolutional LSTMs, the prediction of site α is affected by different types of neighbors (e.g. 8 closest neighbors or 16 second closest neighbors). Thus the hidden cells of site α comprise the hidden cells passed from different types of neighbors and from α itself in the previous state. As illustrated in Fig. 4, we feed all of these hidden cells (denoted by contextual feature maps in Fig. 4) into modified convolutional LSTMs to produce individual hidden cells to α itself and to its different neighbors in the next state. In the following, we will describe the implementation details of diffusion layer. To be self-contained, we first recall the notation of convolutional LSTM [25] and then describe the proposed diffused LSTM.

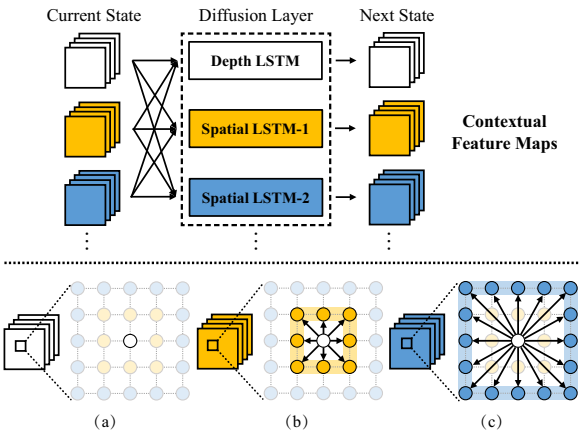

Figure 4. The proposed diffused LSTM layer. **The top row** shows adopting the diffused LSTM to generate contextual feature maps. The input of each LSTM unit is a set of contextual feature maps in current state, while the output is a special type of contextual feature maps for next state. **The bottom row** illustrates the meaning of each type of contextual feature maps. The white maps in (a) are the output of Depth LSTM, and each site indicates the information passed from the corresponding site in previous state. The yellow maps in (b) are generated by Spatial LSTM-1, and each site contains the information that the site will pass to its closest 8 neighbors in the current state. Similarly, each site in the blue maps in (c) denotes the information spreading to its second closest 16 neighbors. Best viewed in color.

3.3.1 Convolutional LSTM

Let $\mathcal{X}_t \in \mathbb{R}^{M\times N\times D}$, $\mathcal{H}_t \in \mathbb{R}^{M\times N\times D}$ and $\mathcal{M}_t \in \mathbb{R}^{M\times N\times D}$ indicate the feature maps, hidden cells and memory cells of state t, where $M\times N$ denotes the spatial dimension of state t and D is the number of channels. Similar to the traditional LSTM [14], we use g^i, g^f, g^c and g^o to indicate input gate, forget gate, memory gate and output gate respectively. And W^i, W^f, W^c and W^o are the gate kernels. Convolutional LSTM accepts the state t as the input and outputs the next state as follows:

$$\begin{aligned} g^i &= \sigma(W_x^i * \mathcal{X}_t + W_h^i * \mathcal{H}_t + b^i) \\ g^f &= \sigma(W_x^f * \mathcal{X}_t + W_h^f * \mathcal{H}_t + b^f) \\ g^c &= \sigma(W_x^c * \mathcal{X}_t + W_h^c * \mathcal{H}_t + b^c) \\ g^o &= \tanh(W_x^o * \mathcal{X}_t + W_h^o * \mathcal{H}_t + b^o) \\ \mathcal{M}_{t+1} &= g^f \odot \mathcal{M}_t + g^i \odot g^c \\ \mathcal{H}_{t+1} &= g^o \odot \tanh(\mathcal{M}_{t+1}) \end{aligned} \tag{1}$$

where $*$ denotes the convolution operator and $\odot$ denotes the Hadamard product. The symbol σ indicates the sigmoid function. Let $\mathbf{W}$ and $\mathbf{B}$ indicate the collection of all gate weights and biases respectively. Following [14] and [19], we define function LSTM$(\cdot)$ to shorten Eqn. (1) with the form,

$$(\mathcal{H}_{t+1}, \mathcal{M}_{t+1}) = \text{LSTM}(\mathcal{X}_t, \mathcal{H}_t, \mathcal{M}_t, \mathbf{W}, \mathbf{B}) \tag{2}$$

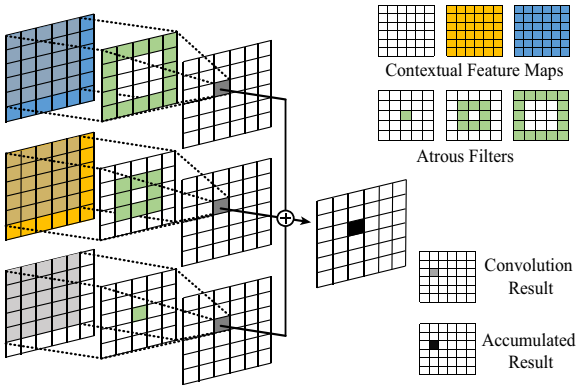

Figure 5. An illustration of convolutional operations in each LSTM unit. A set of contextual feature maps with different meanings (e.g. each site in white, yellow and blue map involve the information pass to self, closest 8 neighbors and second closest 16 neighbors respectively) are fed into LSTM unit. A special type of atrous filters are designed to capture the short range and long range dependencies from various neighbors to a certain site and pass the accumulated information to the next state. For each kernel, the green regions are learnable and others are always set to zero.

The mechanism acts as a memory system, which records the previous information into the memory cells and is used to communicate with subsequent input. In order to ensure the next state has the same spatial dimension as the input, we need to perform zero-padding before applying the convolution operation.

3.3.2 Diffused LSTM

Following the definition in [19], we define two kinds of diffused LSTM in the proposed diffusion layer, named depth LSTM and spatial LSTM. Intuitively, the depth LSTM can maintain the previous information at each site by applying the memory cells benefited from the LSTM mechanism. The spatial LSTM calculates the information that each position travels outward to their neighbors. Note that different spatial LSTM will compute the information that each site passes to the neighbors with different spatial distances. For example, the yellow maps in Fig. 4 are the outputs of spatial LSTM-1, and the value in each position denotes the information that the position propagates to its $3\times3-1=8$ nearest neighbors with distance 1. Analogously, the blue maps are the outputs of spatial LSTM-2, and the value in each position denotes the message passing to the $5\times5-3\times3=16$ second nearest neighbors with distance 2.

As illustrated in Fig. 4, the input image is corresponding to $E+1$ groups of hidden cell maps (we set $E=2$ in this article for illustration), which are generated by one depth LSTM and E spatial LSTM. Let $\mathcal{H}^s_{t,e}\in\mathbb{R}^{M\times N\times D}, e\in\{1,2,...,E\}$ denote the e-th group of hidden cell maps generated from e-th spatial LSTM, and the hidden cells in each position are used to propagate the information to its $(2e+1)^2-(2e-1)^2=8e$ neighbors with distance e. Let $\mathcal{H}^d_t\in\mathbb{R}^{M\times N\times D}$ indicate the hidden cell maps calculated by the depth LSTM using the weights updated in the t-th layer. Thus the gate values of a certain LSTM unit (e.g. depth LSTM or spatial LSTM) in t-th layer can be calculated by,

$$\begin{aligned}
g^i_t &= \sigma(\sum\nolimits_e \{W^s_{t,e}\}^i * \mathcal{H}^s_{t,e} + \{W^d_t\}^i * \mathcal{H}^d_t + b^i_t)\\
g^f_t &= \sigma(\sum\nolimits_e \{W^s_{t,e}\}^f * \mathcal{H}^s_{t,e} + \{W^d_t\}^f * \mathcal{H}^d_t + b^f_t)\\
g^c_t &= \sigma(\sum\nolimits_e \{W^s_{t,e}\}^c * \mathcal{H}^s_{t,e} + \{W^d_t\}^c * \mathcal{H}^d_t + b^c_t)\\
g^o_t &= \tanh(\sum\nolimits_e \{W^s_{t,e}\}^o * \mathcal{H}^s_{t,e} + \{W^d_t\}^o * \mathcal{H}^d_t + b^o_t)
\end{aligned} \tag{3}$$

where $W^s_{t,e}$ and W^d_t indicate the weights of kernels associated with e-th spatial hidden cell maps and depth hidden cell maps in t-th layer. And the superscript i, f, c and o correspond to distinct state gates.

Denote the e-th group of memory cells for the spatial dimension as $\mathcal{M}^s_{t,e}\in\mathbb{R}^{M\times N\times D}$ and the memory cells for depth dimension as $\mathcal{M}^d_t\in\mathbb{R}^{M\times N\times D}$. Same as convolutional LSTM, the novel hidden cell maps and memory cell maps in $t+1$-th layer are computed as,

$$\begin{aligned}
(\mathcal{H}^s_{t+1,1},\mathcal{M}^s_{t+1,1}) &= \text{LSTM}(\mathcal{H}^\dagger_t,\mathcal{H}^d_t,\mathcal{M}^s_{t,1},\{\mathbf{P}\}^s_{t,1})\\
(\mathcal{H}^s_{t+1,2},\mathcal{M}^s_{t+1,2}) &= \text{LSTM}(\mathcal{H}^\dagger_t,\mathcal{H}^d_t,\mathcal{M}^s_{t,2},\{\mathbf{P}\}^s_{t,2})\\
&\dots\\
(\mathcal{H}^s_{t+1,E},\mathcal{M}^s_{t+1,E}) &= \text{LSTM}(\mathcal{H}^\dagger_t,\mathcal{H}^d_t,\mathcal{M}^s_{t,E},\{\mathbf{P}\}^s_{t,E})\\
(\mathcal{H}^d_{t+1},\mathcal{M}^d_{t+1}) &= \text{LSTM}(\mathcal{H}^\dagger_t,\mathcal{H}^d_t,\mathcal{M}^d_t,\{\mathbf{P}\}^d_t)
\end{aligned} \tag{4}$$

where $\mathcal{H}^\dagger_t=\{\mathcal{H}^s_{t,e}\}^E_{e=1}$ is the set of spatial hidden cell maps. $\mathbf{P}=\{\mathbf{W},\mathbf{B}\}$ indicates the parameter set.

In the above process, different numbers of spatial LSTMs allow us arbitrarily enlarge *field-of-view* in the context modeling. For a LSTM unit in the certain layer, there exist $E+1$ filters with distinct forms, and each one is associated with a group of hidden cell maps. As illustrated in Fig. 5, we design a novel type of atrous filters [5] to calculate the convolutional results. Thus depth hidden cell maps specify the depth filter and only the center of the kernel has weight value. The e-th spatial filter is associated with e-th group of spatial hidden cell maps, and it introduces non-zero weights in the sites whose distance to the kernel center is e. Note that, if the site α has neighbor α' with distance e, we need to adopt the hidden cells of site α' in the e-th group to enrich the representation of site α.

In this way, each site in the input image can provide distinct guidance to its neighbors with different distances by employing specific spatial LSTMs, which takes the spatial layouts and interactions into account for feature learning. In order to ensure different neighbors receive various information, the weight matrices $\mathbf{W}^s_t$ and bias $\mathbf{B}^s_t$ of E spatial LSTMs are not shared in this article.

4. Experiments

In this section, we demonstrate the effectiveness of proposed PDNs by comparing it with state-of-the-art semantic segmentation methods. In the following, we first give a brief overview of the datasets and evaluation metrics. Then we report the performance of PDNs on both object segmentation and object part segmentation tasks. The component analysis is also investigated at the end of this section.

Datasets and Evaluation Metrics. We compare proposed PDNs with the state-of-the-art methods on ImageNet Parsing, PASCAL-Person-Part and PASCAL VOC 2012 (VOC2012) datasets. **ImageNet Parsing** [39] is a challenging dataset for scene-centric semantic segmentation task. It includes 150 semantic categories, and many categories have the similar appearance, prompting us to use more reasonable contextual reasoning for dense prediction. In this dataset, 20,210 images are employed for model training and another 2,000 images for validation. **PASCAL-Person-Part** dataset is a fine-grained part segmentaion benchmark collected by Chen *et. al*. [7] from PASCAL VOC 2010 dataset. It contains the detailed part annotation for each person and the annotations are merged into six person parts (i.e. Head, Torso, Upper/Lower Arms and Upper/Lower Legs) and one background category as previous setting in [33, 6, 18]. Totally, 1,716 images are used for model training and 1,817 for test. We also report the performance on the **PASCAL VOC2012** [8] dataset, since it is a standard benchmark for generic semantic segmentation. It contains 20 object categories and the sizes of training, validation and test set are 10582, 1449 and 1456, respectively. For all of the above datasets, we employ standard mean intersection-over-union criterion (denoted as mIoU) and pixel-wise accuracy [22] to evaluate our proposed architecture and other comparative methods.

Implementation Details. In our experiments, our architecture is implemented based on Caffe platform [13] and all networks are trained on four NVIDIA GeForce GTX TITAN X GPUs with 12GB memory. The input image is randomly cropped to 321×321 for model training. Five diffusion layers are stacked onto top of the convolutional neural networks (i.e. ResNet-101*) and we fine-tune the network based on the pre-trained neural network in the bottom. The learning rate of the newly added layers is initialized as 2.5×10^{-3} and that of other previously learned layers is initialized as 2.5×10^{-4}. All the parameters in the diffusion layer are randomly initialized from a Gaussian distribution with the mean 0 and the variance 0.01. We train all the models using stochastic gradient descent with a batch size of 1 image, momentum of 0.9, and weight decay of 0.0005.

*http://liangchiehchen.com/projects/DeepLab.html

Method	pixel acc. %	mean IoU %
FCN [22]	69.05	24.86
DeepLab [4]	71.06	27.08
Grid-LSTM [14]	69.37	25.11
LG-LSTM [19]	69.91	25.79
ResNet [11]	75.09	35.07
ResNet + CRF [11]	77.22	36.79
PDNs	**77.51**	**37.02**

Table 1. Experimental results on ImageNet Parsing *val* set.

4.1. Experiment Results and Comparisons

In the following, we compare the proposed Progressively Diffused Networks with several state-of-the-art methods on image segmentation tasks.

ImageNet Parsing dataset [39]. We compare our model with five state-of-the-art methods: FCN [22], DeepLab [4], Grid-LSTM [14], LG-LSTM [19] and ResNet [11]. These methods can be grouped into two categories: i) CNN-based method: FCN [22], DeepLab [4] and ResNet [11]; ii) LSTM-based method: Grid-LSTM [14], LG-LSTM [19]. The first category employs convolution and pooling operations to directly extract the abstract feature representation. The second category uses LSTM to construct short-distance and long-distance spatial dependencies.

Since ImageNet Parsing is a newly proposed large-scale segmentation dataset, and previous works didn't report their performance on it, all the results in Table 1 are given based on our own implementation. All of the comparison methods have trained 200,000 iterations without any extra data. We record the results every ten thousand iterations and select the best to report. According to Table 1, the proposed PDNs outperforms the baseline method ResNet-101 [11] in terms of pixel accuracy and mean IoU with 2.42% and 1.95% respectively. It well demonstrates that incorporating the contextual information into representation learning can further enhance the discriminative ability of deep features. Fig. 6 gives the visualization results of our method and ResNet-101 [11]. It is obvious that the proposed PDNs can effectively distinguish similar objects in complex scenes by local and global reasoning.

PASCAL-Person-Part dataset [7]. Table 2 shows the comparison results with seven state-of-the-art approaches [4, 33, 6, 14, 19, 18, 11] on the metric of mean IoU. An obvious improvement, i.e. 1.7% increased by PDNs over the ResNet-101, can be observed from the comparison on breakdown categories. In order to reflect the advantages of proposed PDNs in context modeling, we also add CRF [4] on the top of ResNet-101 [11] to optimize prediction results. Such model achieves mean IoU of 62.6%, which is still 1.5% less than our PDNs. The visualization of segmentation results on this dataset are shown in Fig. 7.

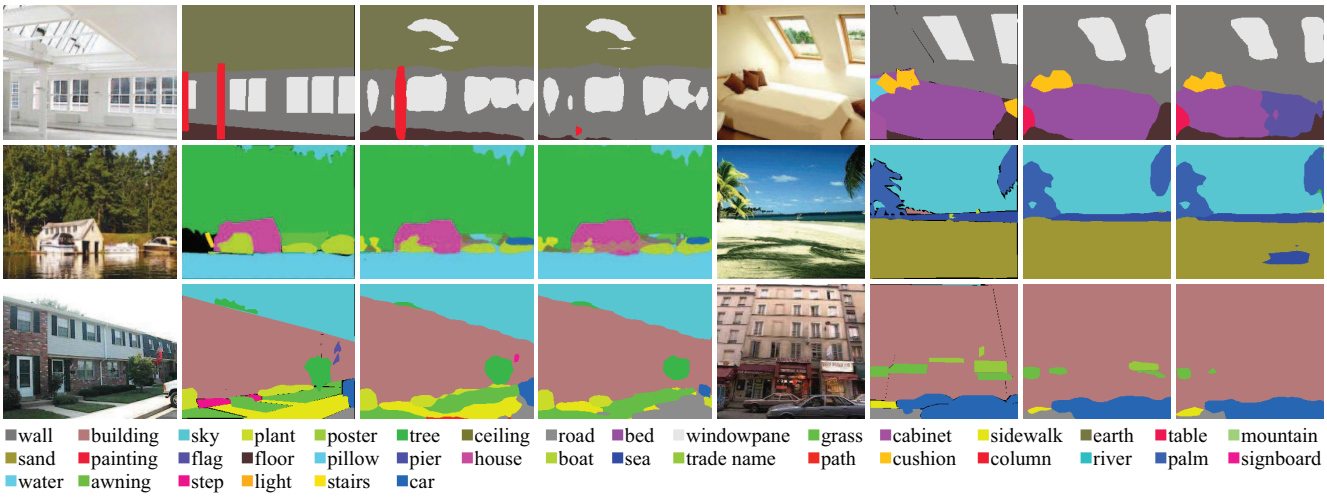

Figure 6. Visualization of segmentation results on ImageNet Parsing dataset [39]. From left to right are input image, groundtruth labeling, segmentation result by proposed PDNs and segmentation by ResNet-101 [11]. Best viewed in color.

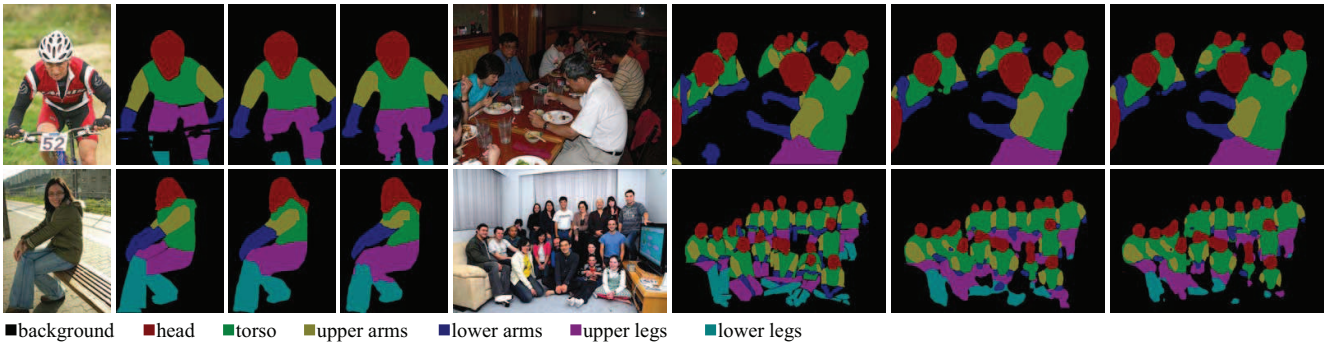

Figure 7. Visualization of segmentation results on PASCAL-Person-Part dataset [7]. From left to right are input image, groundtruth labeling, segmentation result by proposed PDNs and segmentation by ResNet-101 [11]. Best viewed in color.

Figure 8. Visualization of segmentation results on PASCAL VOC2012 dataset [8]. From left to right are input image, groundtruth labeling, segmentation result by proposed PDNs and segmentation by ResNet-101 [11]. Best viewed in color.

Our PDNs gives more consistent segmentation results by incorporating spatial dependencies.

PASCAL VOC 2012 dataset [8]. We report the segmentation results and comparisons with six recent state-of-the-arts [22, 4, 38, 21, 19, 11] in VOC2012. Table 3 shows the breakdown performance of proposed PDNs and comparisons with six state-of-the-arts on VOC2012 *val* dataset. Generally, the proposed PDNs outperforms other methods on vast majority of categories. Compared with the above two datasets, our method doesn't achieve a better improvement than ResNet [11]. One reasonable explanation is that the benefit of context modeling degrades at such dataset. This is because scenes in VOC2012 are less complex than ImageNet Parsing [39] (e.g. less than three foreground ob-

Method	head	torso	u-arms	l-arms	u-legs	l-legs	background	**mIoU**
DeepLab [4]	78.1	54.0	37.3	36.9	33.7	29.6	92.9	51.8
HAZN [33]	80.8	59.1	43.1	42.8	39.0	34.5	93.6	56.1
Attention [6]	-	-	-	-	-	-	-	56.4
Grid-LSTM [14]	81.9	58.9	43.1	46.9	40.1	34.6	86.0	56.0
LG-LSTM [19]	82.7	61.0	45.4	47.8	42.3	38.0	88.6	58.0
Graph-LSTM [18]	82.7	62.7	46.9	47.7	45.7	40.9	94.6	60.2
ResNet [11]	85.1	68.4	50.9	51.2	46.7	39.4	**95.6**	62.4
ResNet + CRF [11]	84.2	68.8	51.2	52.2	46.8	39.4	95.3	62.6
PDNs	**85.6**	**70.0**	**53.3**	**53.2**	**47.5**	**43.7**	**95.6**	**64.1**

Table 2. Experimental results (**IoU**) on PASCAL-Person-Part set.

Method	aero	bike	bird	boat	bottle	bus	car	cat	chair	cow	table	dog	horse	mbike	person	plant	sheep	sofa	train	tv	**mIoU**
FCN [22]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	62.7
DeepLab [4]	78.7	34.3	73.8	61.5	67.2	83.7	78.2	80.9	29.7	74.0	52.1	72.2	68.4	73.6	80.6	48.4	72.4	43.6	77.0	59.1	66.7
CRF-RNN [38]	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	69.6
DPN [21]	84.8	37.5	80.7	66.3	67.5	84.2	76.4	81.5	33.8	65.8	50.4	76.8	67.1	74.9	81.1	48.3	75.9	41.8	76.6	60.4	67.8
LG-LSTM† [19]	77.3	29.7	74.7	58.4	60.6	78.2	76.9	79.2	26.5	51.7	53.6	73.9	62.4	68.1	75.9	43.4	60.7	42.5	78.1	56.7	62.8
ResNet† [11]	**87.9**	41.4	**89.5**	72.5	**80.7**	93.0	87.7	**91.7**	**39.7**	83.2	**53.8**	**85.0**	**85.2**	82.5	85.6	59.8	85.5	40.2	87.0	77.2	76.3
PDNs†-NB24	87.7	42.8	89.4	**73.4**	80.4	93.1	88.8	91.4	39.1	83.7	51.7	84.7	84.8	83.6	86.1	**60.4**	**87.1**	**42.7**	**87.4**	78.2	**76.7**
PDNs†-NB48	87.2	**43.5**	89.2	71.2	78.6	**93.4**	**88.9**	91.1	39.5	**83.8**	47.6	84.8	83.0	**83.9**	**86.9**	60.1	86.8	42.6	86.5	**79.0**	76.4

Table 3. Experimental results (**IoU**) on PASCAL VOC 2012 *val* set. The approaches pre-trained on COCO [20] are marked with †.

jects in an image, obvious differences in objects appearance), and preserving the boundaries of objects is more beneficial than global reasoning. However, in terms of predicting semantic categories with the dramatic changes in appearance and size such as plant, sheep and sofa, our method still achieves a improvement. Fig. 8 gives some exemplar segmentation results on VOC2012.

4.2. Comparison to Different Variants

To strictly evaluate the effectiveness of different numbers of spatial neighbors, we also report the performance of using 24 spatial neighbors and 48 spatial neighbors. We use the capital "NB" to indicate the number of spatial neighboring connections for each pixel. According to Table 3, "PDNs-NB24" with 24 spatial neighbors outperforms that with 48 spatial neighbors by 0.3% over the metric of mIoU. Intuitively, too many neighbors, while bringing more contextual information, also generate much noise, the latter affecting the final prediction results.

We also validate the advantage of using different diffusion layers on the top of ResNet-101 [11] on PASCAL-Person-Part [7] dataset. Table 4 shows the breakdown IoU results, and the improvements can be observed by gradually using more diffusion layers. It verifies well the effectiveness of exploiting more discriminative feature representation by stacking multi-layers to diffuse the information. At the same time, we can find that the performance change between the fourth layer and the fifth layer is small. Intuitively, we can assume that the information has been effectively propagated in the image plane in the fifth layer.

Method	pixel acc. %	mean IoU %
PDNs-1	93.7	63.2
PDNs-2	93.8	63.6
PDNs-3	93.8	63.6
PDNs-4	**93.9**	**64.0**
PDNs-5	93.8	63.9
PDNs-merge	**93.9**	**64.1**

Table 4. Experimental results with different numbers of diffusion layers on PASCAL-Person-Part set.

5. Conclusion

This work proposes a novel progressively diffused networks (PDNs) for context modeling, and pursues the semantic segmentation task. In PDNs, the stacked diffusion layers make the contextual information spreading from a certain site to a large range on the image plane, through adopting convolutional LSTM with special atrous filters to capture the short range and long range spatial dependencies for a certain site and passing the accumulated information to site self or its neighbors in the next layer. The experiment results on three benchmarks demonstrate the effectiveness of proposed architecture.

References

[1] Y. Bengio, A. Courville, and P. Vincent. Representation learning: A review and new perspectives. *IEEE transactions on pattern analysis and machine intelligence*, 35(8):1798–1828, 2013.

[2] W. Byeon, T. M. Breuel, F. Raue, and M. Liwicki. Scene labeling with lstm recurrent neural networks. In *CVPR*, 2015.

[3] G. Y. Chen and G. Núñez. Inflammasomes in intestinal inflammation and cancer. *Gastroenterology*, 141(6):1986–1999, 2011.

[4] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Semantic image segmentation with deep convolutional nets and fully connected crfs. *ICLR*, 2014.

[5] L. Chen, G. Papandreou, I. Kokkinos, K. Murphy, and A. L. Yuille. Deeplab: Semantic image segmentation with deep convolutional nets, atrous convolution, and fully connected crfs. *ArXiv Preprint*, 2016.

[6] L.-C. Chen, Y. Yang, J. Wang, W. Xu, and A. L. Yuille. Attention to scale: Scale-aware semantic image segmentation. *ArXiv Preprint*, 2015.

[7] X. Chen, R. Mottaghi, X. Liu, S. Fidler, R. Urtasun, and A. Yuille. Detect what you can: Detecting and representing objects using holistic models and body parts. In *CVPR*, 2014.

[8] M. Everingham, L. J. V. Gool, C. K. I. Williams, J. M. Winn, and A. Zisserman. The pascal visual object classes (VOC) challenge. *International Journal of Computer Vision*, 88(2):303–338, 2010.

[9] P. F. Felzenszwalb, R. B. Girshick, D. McAllester, and D. Ramanan. Object detection with discriminatively trained part-based models. *IEEE transactions on pattern analysis and machine intelligence*, 32(9):1627–1645, 2010.

[10] A. Graves and J. Schmidhuber. Offline handwriting recognition with multidimensional recurrent neural networks. In *NIPS*, 2008.

[11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CVPR*, 2016.

[12] S. Ji, W. Xu, M. Yang, and K. Yu. 3d convolutional neural networks for human action recognition. *IEEE transactions on pattern analysis and machine intelligence*, 35(1):221–231, 2013.

[13] Y. Jia, E. Shelhamer, J. Donahue, S. Karayev, J. Long, R. B. Girshick, S. Guadarrama, and T. Darrell. Caffe: Convolutional architecture for fast feature embedding. In *ACM Multimedia*, 2014.

[14] N. Kalchbrenner, I. Danihelka, and A. Graves. Grid long short-term memory. *Arxiv Preprint*, 2015.

[15] V. Koltun. Efficient inference in fully connected crfs with gaussian edge potentials. In *NIPS*, 2011.

[16] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.

[17] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural computation*, 1(4):541–551, 1989.

[18] X. Liang, X. Shen, J. Feng, L. Lin, and S. Yan. Semantic object parsing with graph LSTM. In *ECCV*, 2016.

[19] X. Liang, X. Shen, D. Xiang, J. Feng, L. Lin, and S. Yan. Semantic object parsing with local-global long short-term memory. *CVPR*, 2016.

[20] T.-Y. Lin, M. Maire, S. Belongie, J. Hays, P. Perona, D. Ramanan, P. Dollár, and C. L. Zitnick. Microsoft coco: Common objects in context. In *ECCV*, 2014.

[21] Z. Liu, X. Li, P. Luo, C. C. Loy, and X. Tang. Semantic image segmentation via deep parsing network. In *ICCV*, 2015.

[22] J. Long, E. Shelhamer, and T. Darrell. Fully convolutional networks for semantic segmentation. In *CVPR*, 2015.

[23] P. Schnitzspan, M. Fritz, S. Roth, and B. Schiele. Discriminative structure learning of hierarchical representations for object detection. In *CVPR*, 2009.

[24] A. G. Schwing and R. Urtasun. Fully connected deep structured networks. *ArXiv Preprint*, 2015.

[25] X. Shi, Z. Chen, H. Wang, D. Yeung, W. Wong, and W. Woo. Convolutional LSTM network: A machine learning approach for precipitation nowcasting. In *NIPS*, 2015.

[26] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *ICLR*, 2015.

[27] R. K. Srivastava, K. Greff, and J. Schmidhuber. Highway networks. *ArXiv Preprint*, 2015.

[28] I. Sutskever, O. Vinyals, and Q. V. Le. Sequence to sequence learning with neural networks. In *NIPS*, 2014.

[29] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *CVPR*, 2015.

[30] L. Theis and M. Bethge. Generative image modeling using spatial lstms. In *NIPS*, 2015.

[31] A. van den Oord, N. Kalchbrenner, and K. Kavukcuoglu. Pixel recurrent neural networks. In *ICML*, 2016.

[32] V. Vineet, J. Warrell, and P. H. Torr. Filter-based mean-field inference for random fields with higher-order terms and product label-spaces. *International Journal of Computer Vision*, 110(3):290–307, 2014.

[33] F. Xia, P. Wang, L.-C. Chen, and A. L. Yuille. Zoom better to see clearer: Huamn part segmentation with auto zoom net. *ArXiv Preprint*, 2015.

[34] K. Xu, J. Ba, R. Kiros, K. Cho, A. C. Courville, R. Salakhutdinov, R. S. Zemel, and Y. Bengio. Show, attend and tell: Neural image caption generation with visual attention. In *ICML*, 2015.

[35] J. Yang, B. Price, S. Cohen, and M.-H. Yang. Context driven scene parsing with attention to rare classes. In *CVPR*, 2014.

[36] J. Yang, K. Yu, Y. Gong, and T. Huang. Linear spatial pyramid matching using sparse coding for image classification. In *CVPR*, 2009.

[37] L. Yann, B. Yoshua, and H. Geoffrey. Deep learning. *Nature*, 521:436–444, 2015.

[38] S. Zheng, S. Jayasumana, B. Romera-Paredes, V. Vineet, Z. Su, D. Du, C. Huang, and P. H. S. Torr. Conditional random fields as recurrent neural networks. In *ICCV*, 2015.

[39] B. Zhou, H. Zhao, X. Puig, S. Fidler, A. Barriuso, and A. Torralba. Semantic understanding of scenes through the ADE20K dataset. *Arxiv Preprint*, 2016.

[40] Z. Zuo, B. Shuai, G. Wang, X. Liu, X. Wang, B. Wang, and Y. Chen. Learning contextual dependence with convolutional hierarchical recurrent neural networks. *IEEE Trans. Image Processing*, 25(7):2983–2996, 2016.